

## COVER STORY

# Modular Building Blocks for AI on the Edge

By **Frederic Aupetit**, Product Manager, and **Elie Gasnier**, VP of Marketing, ECRIN Systems

Launched in February 2018 at embedded world in Germany, myOPALE building blocks provide truly disruptive technology to realize your design goals for industrial computers that target size, weight, and power consumption (SWaP)-constrained operating environments.

The myOPALE modular concept is based on four key points, all inspired by open or de facto standards used in the embedded market.<sup>1</sup> It eliminates the mechanical link between the CPU module and its I/O cards thanks to PCI Express Gen3-over-cable interconnection. It uses a standard half-height 5.25-inch casing.

Its use of universally deployed Storage Networking Industry Association (SNIA) interconnect standards, such as the Mini-SAS HD (SFF-8643) connector for PCIe-over-cable native protocol, make it directly compliant with new Non-Volatile Memory Express (NVMe) and “just a bunch of drives”/“just a bunch of flash” (JBOD/JBOF) approaches for high-speed network-attached storage (NAS). Last, but not least, it includes the cooling system in each building block.

## myOPALE FOR AI EDGE MODULAR COMPUTING

Based on a COM Express (COMe) carrier board, myOPALE-CPU (**Figure 1**) offers characteristics that extend its application range to environments beyond the limit of standard industrial PCs, including harsher environments such as defense/aeronautics and robotics. ECRIN Systems has qualified its CPU module to withstand shocks of 40 G for 11 milliseconds and 20 G for 20 ms, vibration up to 2 G at 10 to 2,000 Hz, operating temperature of -20°C to 71°C, and relative humidity (RH) of 5% to 90%, providing a cost-effective and robust

alternative to expensive slot-board form factors such as cPCI Serial and VPX in aero/defense programs. I/O extension building blocks can be added to existing myOPALE-IO modules that can integrate one PCIe Gen3 ×8 or two PCIe Gen3 ×4 full-height I/O boards. We are introducing two embedded modules dedicated to artificial intelligence (AI) and EDGE computing.

Today, graphics processors (GPUs) and neural networks create the foundation for leading EDGE computing platforms, bringing


AI and accelerated deep-learning performance as close as possible to industrial internet of things (IIoT) sensors in order to offer real-time processing, low latency, and high accuracy for applications such as smart cities and robotics. ECRIN Systems’ myOPALE-GPU building block (**Figure 2**) allows integration of the embedded Mobile PCI Express Module (MxM) 3.1 GPU mezzanine directly to a Mini-SAS HD PC board adapter, thus eliminating the fragile PCIe PC board connector, a legacy from the IT market that represents a huge point of failure in industrial PCs.

Embedded MxM GPU modules offer low power consumption and the thinnest available commercial off-the-shelf (COTS) solutions for high-performance parallel processing leveraging general-purpose GPUs (GPGPUs). For smart-city AI surveillance, fog computing with augmented reality, widening the field of vision for intelligent logistics, and Industry 4.0 in general, the Nvidia GeForce MxM commercial-grade series (GTX-1030 to GTX-1080 platforms with Pascal architecture and RTX-2060 to RTX-2080 platforms with Turing



**Figure 1:** The myOPALE-CPU building block with its six Mini-SAS HD connectors and 24 PCIe Gen3 links

## Modular Building Blocks for AI on the Edge



**Figure 2:** The myOPALE-GPU building block for AI on the Edge

architecture, offering ray tracing and deep-learning super-sampling technologies) are fine when your requirements do not include long life cycles (18 months only) and big-data processing.

For radar/sonar back-end computers, command/control human-machine interfaces (HMIs) in SWaP-constrained naval environments, test bench systems in aerospace, and medical ultrasound machines, you will prefer a solution offering a five-year life cycle with end-of-life/product change notification (EOL/PCN) process support. Here, the rugged MxM embedded GPUs in the Nvidia Quadro-grade series (Pascal 2000E/GP107 to Pascal 5200/GP104 chip-down) operate from  $-40^{\circ}\text{C}$  to  $71^{\circ}\text{C}$  and offer a new GPU Direct Remote DMA mode that supports the movement of big data for critical missions, with a bandwidth increase of up to 79% (from 3,651 Mbytes/second to 6,556 Mbytes/s) and a latency reduction of 64% (from 41.94  $\mu\text{s}$  to 14.95  $\mu\text{s}$ ).

The myOPALE-GPU offering, which will formally launch in June at the Paris Air Show, has the same general characteristics as the myOPALE-CPU: 5.25-inch box form factor; cold plate or passive heat sink, depending on the environmental conditions; +12-VDC-only power supply (P/S); and PCIe Gen3  $\times 8$  lanes via two Mini-SAS HD connectors. On the front, there are four display ports if the user needs GPU output (not used for GPGPU massive computing). By connecting myOPALE-CPU with myOPALE-GPU, you can build a heterogeneous AI EDGE modular computer composed of a 1 $\times$  Intel Xeon Quad-Core E3-1505L CPU module and 2 $\times$  GPGPU modules via two PCIe Gen3  $\times 8$ . Or you can use 1 $\times$  CPUs with 4 $\times$  GPGPU modules via four PCIe Gen3  $\times 4$  in a very short-depth 2U/19-inch 280-mm chassis or short-depth 1U/19-inch 492-mm chassis.


The second module that will debut at the Paris Air Show is myOPALE-mPCIe (Figure 3), an embedded building block ready to integrate industrial mini-PCIe (mPCIe) I/O mezzanine modules directly to a Mini-SAS HD PCB adapter, again bypassing the fragile PCIe PC board connector.

In this form factor, we find many functions available off the shelf, including communications, networking, wireless, MIL-STD-1553/ARINC avionic bus, and CAN bus. We even have a dual Movidius Myriad X vision processing unit (VPU) with a neural compute engine for enhanced visual intelligence at the network EDGE. It can connect up to 2  $\times$  8 HD resolution RGB cameras directly to the Myriad X VPU, with support for up to 700 million pixels/second of image signal processing throughput. The Myriad X VPU is a power-efficient solution that brings advanced vision and AI applications to devices such as drones, smart cameras, smart homes, security, virtual- and augmented-reality (VR/AR) headsets, and 360 $^{\circ}$  situational awareness.

In many use cases, we will prefer to integrate the AcroPack variant developed by Acromag. AcroPack is an mPCIe-based rugged I/O module for higher density (30  $\times$  70-mm PC board size) to mix and match numerous I/O options in a single building block. Its adherence to the mPCIe standard, a down-facing connector that securely routes I/O signals through the host carrier card without extra cabling, increases reliability. The myOPALE-mPCIe carrier provides direct access to the field I/O signals through front-panel industrial connectors of your choice. AcroPack modules are available as additions to the standard mPCIe functions to support reconfigurable field-programmable (FPGA) functionality, isolated industrial I/O, and octal serial RS. There is also an interesting Gigabit Ethernet (GbE) module with an optional power-over-Ethernet (PoE) function. The AP580E-POE-LF power-sourcing equipment (PSE) device provides 52 VDC at up to 10 watts to a video camera, voice-over-IP phone, or other powered device, letting the user connect point to point without contending with any other traffic on a network. The myOPALE-mPCIe building block, again in the 5.25-inch case form factor, will integrate four or eight mPCIe/AcroPack embedded modules.

Last, but not least, thanks to their use of the SNIA-standard Mini-SAS HD (SFF-8643) connector, myOPALE building blocks are ready for the M.2 form factor and its insanely fast storage interface, NVMe SSD, which uses up to four PCIe Gen3 lanes for a fivefold increase in sequential read/write speed over legacy SSD based on the SATA III interface (to 2,500 Mbytes/s for NVMe, versus 500 Mbytes/s for SATA III). The M.2-to-U.2 (SFF-8639) standard adapter opens the door to myOPALE's use with the many I/O cards and SSD storage options available in this new embedded small form factor. Direct U.2-to-Mini-SAS HD standard cables make it a simple plug-and-play matter to build an embedded video recorder or NAS server based on myOPALE with NVMe storage.

With the cable, you can connect directly to the U.2 connector of a 2.5-inch SSD NVMe or to a rugged NVMe JBOD/JBOF storage system. Trenton Systems proposes a short-depth 2U/19-inch, 410-mm JBOD storage solution with 24 NVMe SSD (equivalent to 192 terabytes of total storage capacity with the 8-Tbyte NVMe SSDs currently available on the market). SSDs are clustered in three independent extra-rugged canisters that hold eight NVMe SSDs each to minimize mean time to repair. Record-breaking bandwidth delivers 27 Gbytes/s of high-performance R/W storage. End-to-end PCIe allows read/writes on up to eight drives simultaneously.



**Figure 3:** The myOPALE-mPCIe building block with its eight rugged I/O sites

## Modular Building Blocks for AI on the Edge



**Figure 4:** Two redundant myOPALE building blocks in a 1U/19-inch rack mount

### FLIGHT TEST AIRCRAFT SYSTEM BASED ON myOPALE

A design win for myOPALE in a flight test aircraft system illustrates the building blocks' capabilities when the system footprint is highly SWaP-constrained. A flight test aircraft embeds a large number of sensors distributed throughout the plane. Various I/O cards connect the sensors to computers. After processing, the test results are stored on removable SSDs for batch processing and replay in the ground station. In our case, more than 3,500 sensors were connected to 14× 1U/19-inch IT servers, integrated in the 14U cabinet height, at a weight of 126 kg with high power consumption. The flexible riser used to integrate PCIe expansion card was not appropriate for the aircraft environment.

To ensure high reliability and a long life cycle, we proposed a myOPALE-Rack of 1U/19-inch short-depth chassis integrating two myOPALE-CPU on the front, two myOPALE-I/O on the back, and two SSDs in one disk tray, accessible from the front (**Figure 4**). Ultimately, the solution was installed in a smaller, 7U/19-inch cabinet, reducing its weight to 67 kg, with a reduced energy cost thanks to the selection of low-power Xeon E3-1515M v5 devices. Each CPU integrated two PCIe slots for expansion cards with a locking system.

### PANEL PC FOR NAVY APPLICATION TAPS myOPALE

Another successful myOPALE implementation shows the benefits of using myOPALE building blocks to create a complete, autonomous system. The solution was for a submarine environment, where the evaluated product life cycle is 30 years. The customer, facing a non-declared EOL, was looking for a solution based on COM Express, as he was convinced that it would be the best solution for equipment sustainability.

The equipment to be replaced was a panel PC consisting of two line-replaceable units (LRUs): the display and the PC box. The PC configuration was quite rich: Xeon CPU, error correction code (ECC) memory, 3× GbE, 4× COM, two PCIe expansion slots for GPU, and two removable SSDs. The thickness of the complete panel PC was limited. Under exceptional environmental conditions, the panel PC had to withstand operating temperatures up to 70°C. On the other hand, the customer needed to validate his software quickly and wanted to have a prototype in a few months.


In this case, the design of a specific COMe carrier was out of the scope of the time-to-market, nonrecurring engineering (NRE) cost, and risk management constraints. To satisfy customer requirements, we proposed an integration based on a myOPALE-CPU and myOPALE-I/O modules. The configuration was accomplished with two M.2 modules operating in a wide temperature range for two extra GbE ports and four RS-232/422/485 serial lines (**Figure 5**).

Without electronics NRE, just a reduced investment for mechanical R&D and thermal qualification of a myOPALE-CPU module equipped with a Xeon E3-1505L v5 (25-W thermal design power [TDP]) in our in-house heating chamber at 85°C for three hours, we met the customer's time-to-market challenge in less than four months. For the customer's software integration and QNX board support package (BSP) adaptation requirements, ECRIN Systems delivered final hardware building blocks and added submodules integrated in a 1U/19-inch rack mount within four weeks ARO.

Your next EDGE computing system based on myOPALE long-life COTS building blocks will be limited only by your imagination. ■

### REFERENCE

1. Aupetit, F. (February 2019). Cable connection concept for remote embedded modular computers. Boards & Solutions. <http://files.iccmedia.com/magazines/basfeb18/basfeb18-p20.pdf>.



**Figure 5:** Navy panel PC based on myOPALE building blocks